Knowledge Representation: How far we have come?

Daniel Khashabi
University of Illinois, Urbana-Champaign

TechReport, Spring 2015

"Concepts are the glue that holds our mental world together"–Gregory Murphy

Abstract

The theory of *Knowledge Representation* has been one of the central focuses in Artificial Intelligence research over the past 60 years which has seen numerous proposals and debates ranging from very philosophical and theoretical level to practical issues. Although a lot of work has been done in this area, we believe currently, it does not receive enough attention from the research community, except in very limited perspectives. The goal of this work is to give a unified overview of a set of select works on the foundations of knowledge representation and exemplar practical advancements.

1 Introduction

Knowledge Representation (KR) is the subfield of Artificial Intelligence (AI) devoted to formalizing information about the world in a form that computer systems can use in order to solve complex tasks, such as a robot moving objects in an environment or answering a question in the form of natural language. The work in knowledge representation goes beyond creating formalism for information, and includes issues like how to acquire, encode and access it, which sometimes is called *knowledge engineering*. Knowledge acquisition is is the process of identifying and capturing the relevant knowledge, according to the representation.

The discussion of knowledge representation has a long history in AI and ranges very fundamental problems like, *whether knowledge representation is necessary for an AI system?* to relatively practical questions like *whether classical logical representation is complete?* Many of the ideas in representation stemmed from observations in psychology about how humans solve problems. Having multiple trends in interpreting the form of understanding, comprehension and reasoning in the human mind has created a diverse set of options for knowledge representation.

In this work we give a relatively comprehensive review on the long-standing works on knowledge representation with emphasis on natural language systems. The review will start from relatively fundamental issues on representation, and will continue to exemplary practical advancements. Due to the close ties between representation and reasoning, various reasoning frameworks have appeared; we briefly explore them in conjunction with representation techniques.

Key terms: Before starting our main conversation we define the terminology we will be using in this summary. We start with proposition.

Proposition: Propositions are *judgments* or *opinions* which can be true or false. A proposition is not necessarily a sentence, although a sentence can express a proposition. For example the sentence that "Cats cannot fly" contains the proposition that cats are not able to fly. Similarly this sentence "Cats are unable to fly" and this declarative phrase with logic `can-fly(cats)=false` convey the same proposition. A claim or a sentence can contain multiple *atomic* propositions. For example "Many think that cats can fly".

Concept: Propositions can be about any physical object (like a *tree, bicycle, etc*) and any abstract idea (like *happiness, thought, betrayal, etc*), which all are (usually) called concepts.

Belief: Belief is an expression of faith and/or trust in a proposition, although a belief might not necessarily be true. For example "Homer believed that the earth is flat" is an proposition which contains the belief of "earth" being "flat", which is not true. "The weather forecast predicts a tornado for tomorrow." is another proposition which expresses another belief which might turn out to be true or false.

Cognition: Discussion of knowledge representation is closely related to the mental picture of the objects and concepts around us. Cognition is the mental action or process of acquiring knowledge and understanding through thought, experience, and the senses. Although cognition involves abstract mental representations of concepts, it can have an outer physical source (like observation, hearing, touching, etc), or inner source (like thinking), or a combination of inner and outer sources.

Knowledge: Knowledge is information, facts, understanding, and skills acquired through experience or education; the theoretical or practical understanding of a subject. The philosophical nature of knowledge has long been studied in *Epistemology* (Steup, 2014), where it has been been classified into different categories based on its nature, and its inherent connection to *belief*, *truth*, *justification*, etc is analyzed.

Representation: To solve many problems there is a need to use some form of information which might be implicit in the problem definition or even not mentioned; therefore this knowledge must be provided to the solver (which can be a computer or any other machine) in order to be used when needed. In other words, representation is a surrogate for a concept, belief, knowledge or propositions. For example, the number 5 could be represented with the string "5", or with bits `101`, or Roman numeral "V", etc. As part of designing a program to solve a problem, we must define how to represent the required knowledge. The theory of Knowledge Representation is concerned with the formalization of how to represent propositions, beliefs, etc. A representation *scheme* defines the form of the knowledge saved in a knowledge base. Any representation is an *abstraction* of the world and its objects. The abstraction level of a representation might be too detailed, or too coarse with only high-level information.

Reasoning: Reasoning is the heart of AI, the decision making system for solving a problem or the action of thinking about something in a logical and sensible way in order to form a conclusion.

Why Knowledge Representation?: It is not a trivial question to answer whether having knowledge representation is necessary or not. In a lot of problems, since dealing with the raw input/output complicates the reasoning stage, historically researchers have preferred to devise intermediate input/output to be middleman between the natural information and the reasoning engine (Figure 1). Therefore the need for intermediate level seems to be essential. In addition, in many problems there is a significant amount of knowledge which either is not mentioned, or it is implicit. Somehow the extra information needs to be provided to the reasoning system, which is usually modeled as a *knowledge base*. The issue of representations applies to both the input level information and the internal knowledge of the reasoning system.

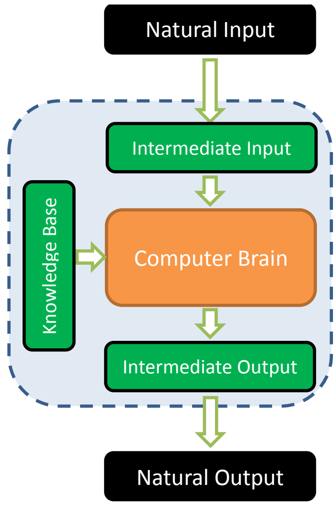

Figure 1: A conventional structure of an AI system. The "intermediate" representation is the middleman between the reasoning engine and the actual input information. In addition to information in the input, for many problems reasoning engine demands knowledge beyond the problem definitions itself, which is provided by a knowledge base. Therefore the discussion of representation can be about either the input information, or the internal knowledge of the reasoning system.

Such separation of knowledge from reasoning system is inherited to us from the very early AI projects and nowadays many accept it as the right way of modeling problems. Whether this is definitely the right way or not, there are debates; at least this is the popular way of doing it. We refer to some of the relevant debates in the forthcoming sections.

2 A brief history of Knowledge Representation

In the early history of AI, there was big series of oppositions to the idea of "abstract representation of thought" among psychology researchers. Perhaps one of the earliest works that created a separate representation of the knowledge from the solver engine was the General Problem Solver (GPS) system of Newell and Simon (1961). This system was intended to be a universal problem solver which could be formalized in *symbolic* form, like chess playing or theorem proving. The idea of symbolic reasoning was developed centuries earlier by Descartes, Pascal, Leibniz and some other pioneers in philosophy of mind. The use of *symbols* as the fundamental elements of the representation, or the *physical symbol system hypothesis* (Newell and Simon, 1961) is a debatable assumption which has continued its way until now.

After the early scalablity issues of GPS, it soon became clear that a system designed for solving any formalized problem is impractical due to the combinatorial explosion of intermediate states for numerous problems. Also, it became apparent early that one of the main problems was how to represent the knowledge needed to solve a problem. Since then, the trend moved towards limiting the scope and focusing on specific problems, and their properties. As a result, the family of symbolic and logical representations, such as propositional and 1st-order logic gained popularity and interesting applications came out (McCarthy and Hayes, 1968). Similarly increased focus on specific applications gave rise to the popularity of *Expert Systems* (i.e. systems which are good on a specific set of problems, rather than everything) (Hayes-Roth et al., 1984). For example, the STUDENT program of Bobrow (1964), written in LISP, could read and solve high school algebra problems which were expressed in natural language; the SHRDLU system of Winograd (1971) with a restricted natural language world model, could discuss and perform tasks in a simulated Blocks World, a famous planning domain with a set of objects (e.g. cubes, etc) on a table where the goal is to build stacks of blocks, given natural language inputs. DENDRAL (Buchanan and Feigenbaum, 1978) was one of the first working expert systems for hypothesis formation and discovery of unknown organic molecules, using its knowledge of chemistry.

There were parallel trend inspired by neurons in the human (Rosenblatt, 1958). This movement started by modeling mental or behavioral phenomena, emergent from interconnected networks of some units. It lost many of its fans after Minsky and Papert (1969) showed fundamental limitations of low layer networks in approximating some functions. However, a series of following events gave another energy to neurally inspired models; Rumelhart et al. (1988a) found a formalized way to train networks with more than one layer. Funahashi (1989) showed the universal approximation property for feedforward networks, i.e. any continuous function on the real numbers can be uniformly approximated by neural networks; Rumelhart et al. (1988b) emphasized on the *parallel* and *distributed* nature of processing, which gave rise to the name "connectionism".

Much progress has been made, but still many problems are unsolved to a large extent. Some issues are far more foundational in AI, like authors who have claimed that human-level reasoning is not achievable via purely computational means; e.g. (Dreyfus, 1992; Searle, 1980; Boden, 1996). Similarly on the representation level, there are many issues that still are subject to debate. For example, *the frame problem* (McCarthy and Hayes, 1968) in logical representation, which is the need for a compact way of expressing states which do not change in an environment as a result of some actions. For example, coloring a table, definitely does not change its position but when using classical logical reasoning, the independence between 'color' property and 'location' property needs to be explicitly encoded. Rather than encoding exceptions in the knowledge base, there has been proposals in the form of non-monotonic reasoning for the frame problem. In nonmonotonic reasoning it is possible to jump to a conclusion and retract some of the conclusions previously made, as further information becomes available (which is the reason for being called *non-monotonic*). Despite the initial promise of the nonmonotonic reasoning methods, many of such reasoning frameworks suffered from the issue of consistency, i.e. ensuring that conclusions drawn are consistent with one another and with the assumptions. Some other works group properties into independent categories to encode causal independence of actions on properties; for example grouping properties with temporal and spatial boundaries (Hayes, 1995).

Minsky and Fillmore were among the first to propose the frame-based representation (Minsky, 1974; Fillmore, 1977), resulted in resources like FrameNet (Baker et al., 1998), or systems like KRL (Bobrow and Winograd, 1977). Frames can be seen as a restricted form of 1st-order knowledge (Fikes and Kehler, 1985). A frame consists of a group of slots and fillers to define a stereotypical object or activity. A slot can contain value such as rules, facts, images, video, procedures or even another frame (Fikes and Kehler, 1985). Frames can be organized hierarchically, where the default values can be inherited the value directly from parent frames.

Frames later on evolved to other representatinal forms many of which are strict subsets of 1st-order logic and some are decidable and closely related to other formalisms such as modal logics (Schild, 1991; Schmidt-Schauß and Smolka, 1991). One of such extensions is Description Logics (also called *concept languages* or *attributive description languages* (Schmidt-Schauß and Smolka, 1991)) that model the declarative part of frames using a logic-based semantics. One of the reasons for such extension was adding more mathematical rigor in hierarchical definition of knowledge. Description Logics (Borgida et al., 1989) guarantee polynomial-time decision algorithms by using operators on concept descriptions (in contrast to the use of quantifiers in 1st-order logic).

What does a good knowledge representation look like? Many researchers have given somewhat different answers. In the following sections we will discuss some of the fundamental issues and research questions regarding representation and related problems. Each section contains a different issue, although there might be overlaps between the issues.

3 Is knowledge representation necessary for AI systems?

The use of a *representation* layer is a common trend among AI researchers. However, (Agre and Chapman, 1987; Brooks, 1990; Brooks, 1991) is among the few works which oppose the common trend. Rather than modeling intelligence via symbol representation, this approach aims to use real-time interaction with the environment to generate viable responses. The big motivation for the design of such *situated* behavior in robots is that most human activity is concept-free, simply *reactions* to changes in their environment (Suchman, 1987; Maes, 1990); for example running, avoiding collisions, etc. Brooks argues that rather than defining knowledge for the robot, it should be able to obtain the behavior via interaction with its environment, since "the world is its own best model". Brooks also argues that human-like intelligence should be evolved via interaction with the environment, rather than defining it directly with representation. In other words, unlike the common top-down approach, "intelligence is determined by the dynamics of interaction with the world". Brook's *subsumption* architecture became successful in applications that needed real-time interaction with a dynamic environment.

Brooks (1991) realizes his ideas with the *Mobots*, which are constructed by linking small state Finite State Machines (FSM). There is an underlying assumption that the Mobot structure can scale up to gain robustness in performance by overlaying more and more specialized mechanisms. Clearly with such a design, the intermediate representations (as the states of FSM's) are inevitable. However Brooks avoids using any direct *declarative representation*. Although this idea found successes in some applications, it did not gain much attention in the majority of AI tasks. Kirsh (1991) and Etzioni (1993) are among the opposing works to Brooks' idea that most human activity is concept-free and AI tasks could be resolved independent of having a direct declarative representation.

First, from a practical perspective, each FSM needs to be tuned to the right stimuli so as to allow the world senses to work properly. This, combined with implicit representation (the state memory of FSM's), can be considered as playing the role of *declarative knowledge*. In addition, the design of FSM's for tasks that demand intricate reasoning about the future is very hard and limited. Therefore one obstacle is the lack of large memory for representation, which seems to be essential for tasks like natural language understanding.

In addition, there are many tasks which cannot be performed based on immediate sensory information alone, but instead require a combination of perception, reasoning, recalling, etc. The decision that a chicken is not a dangerous animal, while an eagle might be, rests in part on the knowledge that chickens are domesticated while eagles are wild. This is despite the fact that there are many similarities between the two birds in terms of sensory input: wings, feathers, beaks, claws. Other tasks may require reasoning or making predictions about the actions of other agents, such as when attempting to dribble past a defender in soccer.

4 Representing Degrees of Beliefs

Is uncertainty with respect to propositions important? Or can we create a complete AI system without taking uncertainty into account? It seems that in the real world, almost any information is subject to uncertainty. Therefore all reasoning problems are pervaded with uncertainty and doubt, which incessantly change along with interactions with the environment. There seem to be many reasons for using probabilistic models for reasoning, although it still has its opponents. Uncertainty may be a result of the uncertain nature of events, disagreement between different facts, or inaccurate or incomplete information. Linguistic imprecision is a very common reason for the occurrence of uncertainty in problems, and unfortunately there is not much direct literature on it, unlike other sources of uncertainty.

How should beliefs with respect to propositions be represented? Is uncertainty deterministically quantifiable, or does it need to be modeled as another level of uncertainty? The very nature of uncertainty has been the subject of epistemological studies for centuries. There are various forms of quantitative uncertainty, i.e. assigning numerical values that express the degree to which we are uncertain about pieces of knowledge. During the eighties there was a big movement to bring in models which support uncertainty for reasoning; for example (Pearl, 1986; Cheeseman, 1985; Zedeh, 1989; Nilsson, 1986; Darwiche and Goldszmidt, 1994; Spohn, 1988). For each of these methods, there are debates over the assumptions, appropriateness and efficiency when used in the reasoning procedure.

Probability Theory is undeniably the most well-known way of representing uncertainty. There is a plethora of works that use probability functions either inside the logical models (Nilsson, 1986), Graphical Models (Pearl, 1986; Lauritzen, 1996), etc. Probability can be viewed as a generalization of classical propositional logic from a binary domain to a bounded continuous domain. There are rigorous studies on agreement between logical representation and probability theory (Cox, 1961; Jaynes, 2003). Despite the popularity of Probability Theory, there are significant limitations and difficulties in applying probabilities, which has resulted in the development of other variations and heuristic techniques for approximating it. There is extensive experimental research in behavioral decision science which verifies that mental judgments often align with probabilistic rules or show an approximation of it. However, there are results which show systematic bias in decision making and strict deviations from norms of probability theory (Tversky and Kahneman, 1974), which is another obstacle in using probability theoretic models. There is much disagreement on the interpretation of probability functions as "degree of belief" (subjective), versus the "long run frequency" (frequentist or objective) interpretation, etc (Hjek, 2012). Bayesianism which currently plays a big role in probability theory, uses the subjective interpretation by using priors as an amount of belief in a proposition combined with the new evidence, which results into a posterior probability.

Dempster-Shafer theory (DS) (Shafer, 1976) can be seen as a generalization of the Bayesian theory. It represents uncertainty with belief functions, which are a way of representing epistemic plausibilities, and are not necessarily the same as probability functions. DS also provides a way of combining beliefs. From a computational point of view, higher-order representations such as Dempster-Shafer are harder than the probabilistic models.

Fuzzy Logic (Zadeh, 1973) can be thought of as the generalization of traditional logic to contain a truth value that ranges in degree between 0 and 1. Fuzzy logic deals with cases where something could be partially true; we can think of it as dealing with "shades of grey" rather than the "black or white" of classical logic.

5 Dealing with commonsense

Commonsense is usually defined as "knowledge about the everyday world that is possessed by all people" (Liu and Singh, 2004). It is usually the set of information which is possessed by most people in a society; e.g. objects fall on the ground; chickens do not fly; friends support each other; Barack Obama is the US president [1].

Commonsense knowledge is gained, mostly, during human experience and interaction through our sensory signals about spatial, physical, social, temporal, and psychological aspects of everyday life. Since the acquisition of such knowledge is done automatically during our daily life, it is assumed that every person possesses commonsense, and therefore typically it is never specified in any sort of communication and text.

In many problems in natural language understanding, there is a need for a surprising amount of commonsense (for example, question answering on elementary-school stories (Richardson et al., 2013), or Textual Entailment (Dagan et al., 2010)). Commonsense can play an important role, even in mature NLP problems; for example take the task of parsing the phrase "animals other than dogs such as cats", which is relatively easy for human. For a computer parser it can be ambiguous whether "cats are animals" or "cats are dogs" is more plausible, although it is clear to humans that cats cannot be dogs, based on our commonsense (Wu et al., 2012).

Early AI, during the sixties and onward, experienced a lot of interest in modeling commonsense knowledge. McCarthy, one of the founders of AI, believed in formal logic as a solution to common sense reasoning (McCarthy and Lifschitz, 1990). Marvin Minsky, in his famous book, estimated that "... commonsense is knowing maybe 30 or 60 million things about the world and having them represented so that when something happens, you can make analogies with others" (Minsky, 1988). There are works focusing on elementary-school level story comprehension, which is full of commonsense facts (for example, Charniak (1972) and Dejong (1979)). Unfortunately many of such efforts were tested on very limited problems, like the Block World problem. There have been famous decade-long efforts to create knowledge bases which contain commonsense, such as Cyc (Lenat, 1995) and ConceptNet (Liu and Singh, 2004) (explained in the last section), but none of these have been shown to be a good solution to NLP problems, at least to a reasonable extent.

Starting in the nineties, most of the excitement about commonsense disappeared and the focus shifted towards side problems which seemed to be easier and approximations to bigger challenges. Unfortunately there are not many success stories in literature on this issue.

Recently there have been multiple initiatives to attack problems which need deeper semantic reasoning (such as commonsense reasoning), e.g. MCTest QA challenge (Richardson et al., 2013), Winograd Challenge (Levesque et al., 2011), Facebook's QA dataset (Weston et al., 2015), high school science test challenge (Clark, 2015); although no reasonable solution yet exists for any of the aforementioned challenges. One common property of these challenges is that all of them need extensive knowledge which is not directly mentioned in the problem or is at most implicitly mentioned, which makes it essential to have a knowledge layer along with reasoning, and stresses the importance of the representation problem.

6 Explicit representation vs. distributed representation

In this section we provide a couple of major arguments for and against connectionsm and classifical formalisim (such as logical modeling). Before jumping into the arguments, we find it important to mention that making a seamless comparison between connectionism and classicism is almost impossible, since they are based on different assumptions about the availability of resources (such as data and computation) or the task being solved. Thus one needs to be very careful in ac-

[1] This is commonsense for American people, but not necessarily in another country.

cepting arguments for or against these modeling frameworks.

How is knowledge encoded in human mind? Even after decades of study of the human mind, there is still debate on to what extent knowledge is localized in mind. Some argue that all knowledge is spread across the entire cortex and patterns of neural firing responsible for representing some external event. On the other hand, some believe that certain kinds of information are stored in tightly circumscribed regions. Since the structure of brain has been a source of inspiration for AI researchers, this issue has found mathematical elaboration, giving rise to models and philosophical issues. Connectionism is based on the assumption that cognition takes place via interactions between large numbers of simple processing units, linked via weighted connections. Each concept is represented by many neurons, the weights of the connections and the topology of the network (Touretzky and Hinton, 1985). Similarly, each neuron participates in the representation of many concepts. This is in contrast to 'localist' representation which uses one symbol per concept. Localist designs are easy to comprehend and design, but inefficient when data has big componential structure.

The first few decades of AI were dominated by symbolic representation. Symbolic logic provides well-understood declarative knowledge representation and reasoning paradigms. One big problem with these models is the need for a huge number of knowledge definitions. Otherwise, there will be a sparsity issue and a lack of robustness as a result of missing information in knowledge. Distributed representations tend to be more robust[2], since their models are based on vector representation of input information (say, words). However, unlike logical representations, they lack reasoning like induction, deduction, etc. Instead they just learn a mapping from input to output space. Learning such a mapping usually needs a large number of training examples in order to exhibit enough generalization. Conversely, classical models are capable of learning immediately upon observing the input in formal form, without extra repetitions.[3] Another issue is the degradation of previously learned information in connectionist models, upon learning new facts (also known as known as *catastrophic interference*). While in logical models, when adding more axioms to the system, with the use of a sensible reasoning system, there is no loss of previous information.

In many AI designs, a considerable amount of initial knowledge is given to a learning agent. There has been a lot of debate about whether humans are born with knowledge or not. Pinker and Bloom (1990) and Chomsky (1988) argue that children are born with innate domain-specific knowledge of the principles of grammar, which is, perhaps a support for classical AI systems, in which the system is given expert knowledge while modeling. However, Elman (1996) argues against the innateness hypothesis, which is a support for models learning from scratch, similar to learning statistical models with random initialization (including many connectionist models).

One of the arguments against the connectionism models was proposed by (Fodor and Pylyshyn, 1988), where they questioned the ability of connectionist networks to embody systematicity and compositionality. Fodor and Pylyshyn argue that any reasonable model of representation needs to have these properties; "Compositionality" is the property of natural language whereby one can create novel meaningful elements from smaller constituents. "Systematicity" is the latent connections created in language to link certain roles and constituents of sentences to create meaningful sentences. For example, for any English-speaker the sentence "Jack loves Kate" makes sense, although it might not make much sense if we randomly scramble the words. But we understand that "Kate loves Jack" is probably implied from the previous sentence. Such observations are the basis for the *sub-symbolic* hypothesis that the brain must contain symbolic representations similar to the constituents of language (although independent of the language being spoken, hence called the "language of thought"). How can such linguistic phenomena which seems to have a discrete symbolic

[2]This phrase is a little inaccurate. A connectionist model needs to see many learning instances in order to generalize enough, just like the need for many rules in the classical frameworks.

[3]For human mind, this is subject of some debate. Some believe that 'rehearsal' is a necessity before saving memories into long-term memory (cf. Goldstein (2014)).

nature be analyzed with a network of functions? The criticism of Fodor and Pylyshyn (1988) was followed by a series of responses arguing that systematicity and compositionality can be simulated by means of complex networks (e.g. work by Smolensky (1988)). Elman (1991) proposed a solution by introducing recursive neural networks which repeatedly uses it's previous output and a part of the context to model the recursive behavior of the language.

There have been many studies on the structural behaviors inside learned networks following empirical behavior observations in human. For example "U-shaped" development of cognition has been observed in human when learning various tasks. Usually when we are learning an action (say the past form of verbs), the performance degrades as a result of over-generalization of exceptions. This is followed by an increase in performance by learning the right level of generalization of rules. Rumelhart and McClelland (1985) observed such behavior in their network when learning past tense for English verbs.

Another issue is the "binding" of variables, which has a close connection to the "systematicity" issue raised by Fodor and Pylyshyn. Binding is the act of representing conjunctions of properties. A property can be anything being sensed, for example: color, shape, orientation, etc. To visually detect a red circle, among a blue circle and a blue rectangle, one must visually bind each object's color to its shape (Treisman and Gelade, 1980). Similarly, one can imagine *thematic roles* for groups of words in sentences which binds them together (Fodor, 1983). For example, in order to understand the statement, "Jack feels Kate is angry" one must bind "Jack" to the agent role of "feels". Binding is the backbone of symbolic representation. The real world is full of concepts with complex combinations of properties which require efficient dynamic binding mechanisms in order to generate representations of perceptual objects and movements. However it is not obvious how to do dynamic binding in a mathematical framework. One of the popular trends following neuroscientists' observations of the human body (Gray and Singer, 1989), is the binding base on *temporal synchrony* of the connected units: if two units are bound, then they fire synchronously; otherwise they fire asynchronously.

One big practical limitation of many proposed ideas for dynamic binding is that they are inefficient when it comes to modeling facts and long-term memory. Smolensky (1990) proposed tensor product for simulating the process of variable binding, where symbolic information is stored at and retrieved from known locations. One of the frameworks for natural language processing created based upon temporal synchrony was SHRUTI (Shastri and Ajjanagadde, 1993), in which dynamic bindings are modeled by synchronous firing of appropriate role and entity cells. The first-oreder logic type knowledge is modeled as directed connections. Inference in SHRUTI corresponds to a transient propagation of rhythmic signals over the network. In practice creating such a network might require a gigantic space of nodes and links which is able to process the information at each node in parallel. The authors at the time were motivated by the empirical results suggesting the possibility of synthesizing such structures.

The debate about systenaticity instigated a series of works on neural-symbolic models which aim at the integration of neural networks and symbolic knowledge. There might seem to be a rivalry between connectionism and symbolic representation, however some believe that the way the mind works is a combination of both. Although the mind implements a neural net, it is also a symbolic processor at a higher level. When we do many daily routine skills, such as opening a door or riding a bicycle, it seems like we do not think about details of the actions; we never reason about how to move our weight so that we keep our balance. Instead we do the actions unconsciously and without much reasoning; overthinking the details of our actions might actually confuse us. This perhaps might be an evidence for implicit knowledge in mind. On the other hand, the first time we want to drive a manual car, we might need to get instructions from someone, follow instructions of a book, or observe someone who knows how to do it (each of which being different approximations to what actually needs to be done). This part of our learning is more similar to a classical form of learning (e.g. learning from declarative/procedural definitions). When learning, while repeating the actions we need to carefully think about them and the order in which they must be done. By repetition of actions, they become part of our skill, i.e. we do not need much thinking/reasoning while doing it. This can perhaps be seen as transfer of knowledge/skills from explicit form to implicit representation.

Following similar intuitions, a series of works appeared under "Hybrid Networks" or "Neural-Symbolic Networks". For example Towell and Shavlik (1993) proposed KBANN (Knowledge-Based Artificial Neural Network), a system for insertion and refinement of a neural network with backpropagation, which also has the ability to extract rules from the resulting network. A pure connectionist model can be viewed as learning a set of propositional rules. Minimizing the energy function of a symmetric network corresponds to finding a model compatible with a set of propositonal rules (Pinkas, 1991). But there is hope to gain capabilities beyond propositional logic in connectionist models. Ballard and Hayes (1986) formulated resolution on a restricted family of 1st-order rules as an energy minimization problem on a network. Hölldobler et al. (1991) developed a method for encoding propositionl logic inside a multi-layer feed-forward network. It has been claimed that they can create a loop between the continuous network and the symbolic logic by extracting the logical rules from the network and feeding them into the network.

7 The issue of abstraction level

One of the issues which is arises as a result of explicit representation of facts and rules is the "abstraction" issue, which unfortunately does not have much direct literature. Making rules more coarser sometimes makes processing and representation easier. For example suppose in a database of people the final task involves whether people are engineers or not. Consider these two scenarios: having a database of people with (1) with more abstracted form of job title, i.e. "engineer" or "non-engineer", or (2) full job names, such as "psychologist", "mathematician", "chemical engineer", etc. In case (1) the decision is relatively easier that case (2), in which the decision needs further collapsing of the job-names into "engineer" and "non-engineer". Therefore, usually the more detailed the information is, the more computationally difficult it is to reason with. In general there is a trade-off between the expressive level of the representation and the deductive complexity.

What if in the previous example, the goal is a decision dependent on the exact job name? Then case (2) would be the right abstraction level and case (1) a bad approximation to the information needed. This shows that the right level of abstraction depends on what is needed. The abstraction level should be in a way that it is "expressive enough" to find the answer, without further additional complications.

Since many levels of abstraction might be necessary for reasoning, it might be a good idea to model an environment at multiple levels of abstraction (Rasmussen, 1985; Bisantz and Vicente, 1994). Such hierarchical abstract can be modeled in various ways, 1st-order logic, frames, semantic graphs, etc, and the properties are inherited from parent nodes to more refined nodes (for example work with frames (Bobrow and Winograd, 1977)). One issue with the default pass of properties to children is how model exceptions. For example, if any "bird" can "fly", any entity of type "bird" will inherit the "fly" property. What is the right way handle the exception that "chicken" does not fly?

Another issue is during accessing the knowledge. Suppose a problem is given to the system, and internally it needs to choose which level of abstraction to use. What is the right way of determining what is the right level of abstraction? For example, two different inputs to a questions answering system might be (1) "Is there any person with degree in psychology?", or (2) "Is there any person with degree in a non-engineering major?". The attention structure in human mind (Johnson and Proctor, 2004; Janzen and Vicente, 1997) is very strong in finding out what are the right abstraction for answering each question, but for computers, it might need a further pre-processing on the question to decide where to look for information. The issue of accessing knowledge goes beyond abstraction levels; in general it can be about accessing multiple knowledge bases, possibly with contradictions, etc. Unfortunately there is not much formalism on these issues.

8 Exemplar Practical Systems

Here we give a summary of the main efforts in creating knowledge bases. Many of the following KBs claim to to capture general-purpose world-semantic knowledge; however the small differences in their knowledge representations make them suitable for very different purposes. In each of the following we will give partial answers to the following questions:

- What are the representations? What is the level of abstraction?
- How is knowledge acquired?
- How are contradictions and updating the existing information handled?
- Is there any notion of time/date/number? If so, how are temporal and quantitative issues handled?
- How the information is accessed?
- Prominent success stories.

The works are presented in chronological order.

8.1 Cyc (1984–present)

The Cyc project (Lenat, 1995) started with the effort to formalize knowledge (including commonsense) into a logical framework. As of 2003, Cyc contained around 118,000 concepts, used in around 2,000,000 assertions (Liu and Singh, 2004). The assertions in this KB are mostly crafted by knowledge engineers.

To properly use Cyc KB it is essential to first map the raw text into CycL, the logical representation used by Cyc. However, the mapping is quite complex due to the complexity and ambiguity of natural language. Since the representation of concepts and assertions are formalized based on logic, its deductive reasoning works great when all details are precise and unambiguous. However, the usual nature of natural language is full of ambiguity and brevity. The limited access to the knowledge base and tools to use it for common public was another limitation which made it hard for researchers to improve upon this KB.

8.2 WordNet (1985–present)

WordNet (Miller, 1995; Fellbaum, 1998) is arguably the most impactful resource used in NLP community. The resource is optimized for lexical categorization. It contains (mostly) nouns, verbs, adjectives and adverbs[4], labeled with 'sense' classes. Meaning in Wordnet is defined through synsets, which are the set of synonyms. In addition, the elements are linked with small set of semantic relations, e.g. 'is-a', synonym, hyponym, etc. For example here, the hierarchy of hypernyms of `happiness` is shown bellow:

```
abstraction
  => attribute
    => state
      => feeling
        => emotion
          => spirit
            => emotional_state
              => happiness
```

In other words, WordNet could be seen as a combination of dictionary and thesaurus. It also contains the morphological information to get the lemma or stem of a word. As of 2012, it contains 155,287 words with a total of 206,941 word-sense pairs, organized in 117,659 synsets.

The knowledge in the WordNet is mostly handcrafted by knowledge engineers. WordNet has extensively been used for word sense disambiguation (WSD), in which it is aimed to differentiate the right meaning of the words, based on context information. Given the nice relational hierarchy of words in WordNet, one of its popular usages is as a similarity metric, for example by taking into account the distance in the hypernymy tree, or similarity between the synsets.

A criticism (or maybe advantage?), is that the information in WordNet is general, and there is no domain-specific classification of knowledge (e.g. food, engineering, etc). Also it has been argued that the sense categorization inside WordNet is too fine-grained. For example, it can happen in many cases that, a word's meaning corresponds to multiple senses in WordNet. To solve the sense granularity issue, there have been many proposals, including clustering techniques, etc.

[4] It ignores determiners, prepositions, pronouns, conjunctions, and particles.

8.3 ThoughtTreasure (1994–2000)

Begun in December 1993, ThoughtTreasure (Mueller, 1998) was aimed to contain commonsense knowledge. Adopting the view of Minsky (1988) that there is no single "right" representation for everything, ThoughtTreasure uses various forms of representation to represent the knowledge. The major part of the knowledge is a set of *concepts*, which are linked as *assertions*. Bellow are examples of the assertions:

```
[isa soda drink]
(Soda is a drink.)


[is the-sky blue]
(The sky is blue.)
```

The final version contained a total of 27,000 concepts and 51,000 assertions. Unlike many other KBs, it has several domain-specific lower ontologies such as for clothing, food, and music. It contains assertions containing quantities, such as the followings:

```
[duration attend-play
NUMBER:second:10800]
(The duration of a play is 10,800
seconds.)

@19770120:19810120|[President-of
country-USA Jimmy-Carter]
(Jimmy Carter was the President of
the USA from January 20, 1977 to
January 20, 1981.)
```

The ontology supports English and French; 35,000 lexical entries in English and 21,000 lexical entries in French. For each entry there are at most 118 features attached. Examples of features are `ZEROART` (zero article taker), `SING` (singular), `FML` (formal), `CAN` (Canadian).

Grids are 2D maps of floorplans which represent in a typical life how objects (flower, window, table, chair, etc) are arranged near each other. A grid can belong to a kitchen, bedroom, a restaurant, etc. There are rules for how the grids (for example bedroom and kitchen) might be related to each other in the form of procedural commands. There are 29 grids in total included in the KB.

ThoughtTreasure also contains procedural knowledge, in the form of about 100 *scripts*, which are representations of typical activities (e.g. eating at a restaurant). For a fixed script, it contains details details on how events follow each other, and how people and physical objects might interact with each other. For example the procedure to buy a ticket, is a represented as an automaton:

```
purchase-ticket(A, P) :-
    dress(A, purchase-ticket),
    RETRIEVE building-of(P, BLDG);
      near-reachable(A, BLDG),
    near-reachable(A, FINDO(office)),
2:  interjection-of-greeting(A, B =
    FIND(human NEAR counter)),
    WAIT FOR may-I-help-you(B, A)
      OR WAIT 10 seconds AND GOTO 2,
      ...
```

The KB contains such scenarios for typical activities like inter-personal relations, sleeping, attending events, sending a message with phone, etc. The finite automata give a relatively good way to simulate behavior of actors in different scenarios.

As mentioned, for different types of information different ways of representation have been used; (1) logical assertions for encyclopedic facts, (2) grids for stereotypical physical settings, and (3) finite automata for rules of thumb, device behavior, and mental processes.

Considering the lexical hierarchies by some specific set of relations, ThoughtTreasure is similar to WordNet. Also it is similar to Cyc, in the sense that it has commonsense knowledge and most of its information is as assertions between concepts. However, as mentioned earlier, ThoughtTreasure is using multiple representations in addition to logic (finite automata, grids, scripts).

8.4 ConceptNet(2000–present)

ConceptNet (Liu and Singh, 2004) was motivated by importance of *commonsense* knowledge in textual reasoning.

As of 2014, ConceptNet5 (Speer and Havasi, 2012) contains over 10 million assertions, and about 7 million of those assertions are in English. There are 50 relations, and 2,798,486 English concepts. As a graph, just the English part is about 3 million nodes and 7 millon edges or arcs connecting them.

Compared to WordNet, it contains more sophisticated relations like *LocationOf*, *SubeventOf*, *UsedFor*, etc. An example subgraph of the conceptNet is shown in Figure 2.

The data collection for ConceptNet was extracted from the Open-Mind Commonsense (OMCS) project (Singh et al., 2002), during which common people (rather than knowledge experts) were asked to fill in possible relations between concepts or a concept relating to a another one. The OMCS data was already in a semi-structured form the way it was collected by prompting users with fill-in-the-blank templates. The extraction of ConceptNet5 relations are done by design of regular expression on the semi-structured OMCS data. In the final stage of the construction, the concepts are stripped of determiners, and tenses and reduced to their *lemma* forms.

ConceptNet's knowledge representation is semi-structured English. Since many of the concepts are represented with natural language, there are many redundancies in the representation. Unlike WordNet, the concepts can be relatively more complex phrases, e.g. 'cut food', 'kick a ball', 'drive to work'. Also there is relativey more variety of relations between concepts; e.g. 'IsA', 'RelatedTo', 'PartOf', 'UsedFor', 'CapableOf', 'AtLocation', 'Causes', 'Synonym', 'Antonym', etc. Also, conceptNet, does not have sense labels (either relations or concepts). For example the concept 'Power' (in Figure 2) can be used in the sense of 'having authority', or in the sense of 'rate of doing physical work'.

8.5 OpenIE (2003–present)

Open Information Extraction (Etzioni et al., 2004; Mausam et al., 2012) consists of a series of projects dedicated to extraction of relational tuples from text, without resorting to any pre-specified vocabulary or ontology. The systems mostly start with identification of relation phrases and associated arguments, for a given arbitrary sentences. The type of the relation can be very general; mediated by a verb, nouns, adjectives, etc.

In the following box, a sample output of OLLIE (Mausam et al., 2012) on the input text has been shown.

Q: Jimbo was afraid of Bobbert because she gets scared around new people.

```
(she;gets scared around;new people)
(Jimbo; was afraid of; Bobbert)
(Jimbo; was; afraid of Bobbert)
```

Although OpenIE, at first look, is an open relation extractor, the result of running it on web-scale data can be used as a knowledge base of relations between arguments. In Balasubramanian et al. (2013) unigram and bigram counts of co-occurring relations have been published.

Unlike other knowledge bases, OpenIE is an open relation-tuple extraction system which can be used a resource. Therefore, there is not much human power is used for extraction of the rules. One issue is that, since the rules are extracted from text, it will have problems in extraction of the rules that are implicit or not usually mentioned (like facts relevant to commonsense). It has also been argued that the surface relations for OpenIE are brittle. In other words, a fixed relation might be expressed in many different forms, and they will be treated differently. Similarly, a surface string, might have two different meanings, depending on its context, therefore expressing two different sense of the relation.

8.6 Freebase (2007–2015)

Freebase (Bollacker et al., 2008) is a collaborative knowledge base acquired by Google in 2010. Including 46 million of concepts, Freebase is a very popular source of information. The information was acquired from many sources, such as Wikipedia page and human entries. The data are represented in the form of entities/concepts which are connected by relations. The KB can be searched, queried like a database and used to provide information to applications. In 2015, Google announced that it is shutting down Freebase, and will be replaced with Wikidata (see Section 8.9).

There are many successful usages of Freebase, although many of such works are limited in the sense that, they mostly rely on direct usages of the facts, without much reasoning on chains of facts.

8.7 NELL (2010–present)

The Never-Ending Language Learner (NELL) (Carlson et al., 2010; Mitchell et al., 2015)

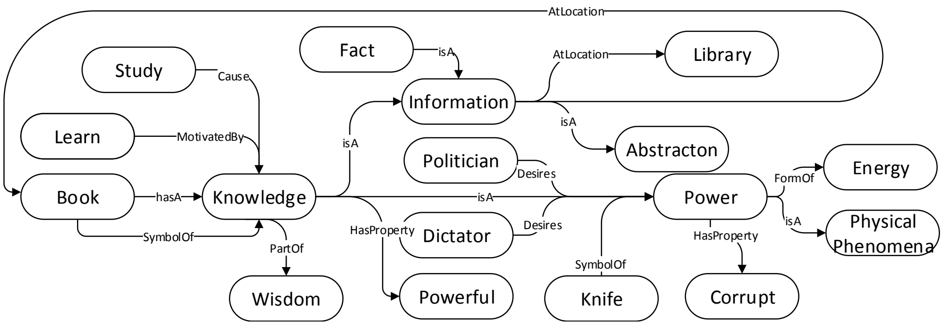

Figure 2: A small visualization of the relations inside ConceptNet5. From `conceptnet5.media.mit.edu`

has been developed to learn from the web 24 hours/day since its start in January 2010. So far has acquired over 80 million beliefs (relations among two fixed entities/concepts). NELL updates its parameters based on new belief, which results in extension of its capability for generating new beliefs.

NELL is one of the most unique efforts to generate knowledge in the loop of reasoning and usage. There has been much engineering on this system to create a stable loop of reasoning and knowledge extraction. For example, deviations in knowledge base facts might be result of some initial erroneous conclusions which creates systematic noise in the results, as the loop of reasoning and extraction continues based on the previous facts. There is vast number of new innovative works suggesting ways to add new facts given the existing knowledge; e.g. (Gardner et al., 2013; Lao et al., 2011). Despite the huge volume of facts in NELL, the usage of this KB has been limited in problems which need reasoning on top of facts, e.g. question answering.

8.8 Probase (2012–present)

Probase (Wu et al., 2012) is a probabilistic taxonomy which contains 2.7 million concepts extracted automatically from 1.68 billion web pages. The probabilities are used to model inconsistent, ambiguous and uncertain information.

As mentioned Probase is a *taxonomy* which can be thought of as an ontology with only *isA* relations. The extraction of concepts is done automatically which has resulted in its big size. The extraction is done by bootstrapping using syntactic patterns and using iterative use of semantic information in the extracted patterns. Compared to Cyc, which has been using knowledge experts for many years and Freebase, which was using communal effort it can be considered an advantage. From this perspective it is similar to OpenIE, NELL, etc. Unfortunately Probase has so far been used only for Microsoft's internal user and not publicly available.

8.9 Wikidata(2012–present)

Wikidata (Vrandečić and Krötzsch, 2014) is a collaboratively knowledge base operated by the Wikimedia Foundation.

The data is distributed across documents which represents a topic and is identified by a unique number (for example the item for the topic Politics is Q7163). Each topic contains a series of statements which are in the form of key-value pairs. Each single user can edit the values or add new keys. The current user-interface of Wikidata seems to be user-friendlier Freebase, another collaborative KB. The creators believe that such the community-driven design will be the way to create an ever-evolving and robust knowledge base.

8.10 Other efforts

Two early frame-based systems, FRL (Roberts and Goldstein, 1977) and KRL (Bobrow and Winograd, 1977) both defined on based on frame models with default assignments, are the early exemplar usages of frame-based knowledge bases, although tailored for limited problems. MindPixel(2000–2005) was a web-based collaborative project aimed to create a KB, though validating true/false statements or probabilistic propositions by online users. Yago (2007–present)

(Suchanek et al., 2007) and DBPedia (2007–present) (Auer et al., 2007) are two other knowledge bases which are similar in representation to Freebase. Their data is mostly collected based on Wikipedia data. BabelNet(2012–present) (Navigli and Ponzetto, 2010) is both a multilingual knowledge base, with a semantic network which connects concepts and named entities in a very large network of semantic relations. It contains about 14 million entries, in about 271 languages.

It is important to mention that, there has been a surge of interest in creating distributed representations for natural language; for example `Word2Vec` (Mikolov et al., 2013; Pennington et al., 2014) and many other ongoing works.

9 Concluding remarks

Here we reviewed a substantial range of works related to Knowledge Representation. However due to the richness of this subarea, much is missing here. Our goal was to give an overall view of a series of outstanding past works to the readers who might not have had the chance to go over this vast literature, and remind ourselves the challenges the modern trends in AI and Machine Learning are facing.

Since the advent of AI systems, many works have been based upon the *physical symbol system* (Newell and Simon, 1976) assumption. This seems to be a strong assumption, but just like many scientific hypotheses, it is made by empirical observation and it might not be necessary (or might even be wrong).

Another trend followed traditionally is the separation of knowledge bases from reasoning systems, which we have inherited from the first systems designed based on logical axioms of world, along with an independent logical reasoning (McCarthy, 1963).

There is not much work directly addressing the issue of *knowledge access*. In real problems when there are multiple levels of abstraction in knowledge bases, or multiple resource, the question becomes what is the right place to look for knowledge, and how to use it. This seems to be an open issue to a large extent.

It seems to be a prevalent idea that uncertainty is a necessity for AI. Despite decades of working on probability theory and the whole set of areas generated based on it, probability might not be the right solution for simulating uncertainty in AI, although there is no doubt that it is a useful solution.

Although we did not focus on reasoning, it is directly determined by the way information about the problem are formalized. Defining what is the reasoning should be is hard; deduction, induction, family of non-monotonic reasoning, MAP, etc? Each of these have their own properties which others might not have, but is there any *efficient* reasoning which subsumes others? Certainly if there exist any, it needs to handle the diverse set of representations too.

Acknowledgment

I would like to appreciate Dan Roth for his invaluable advice and tremendous discussions when writing this report.

References

Philip E Agre and David Chapman. 1987. Pengi: An implementation of a theory of activity. In *AAAI*, volume 87, pages 286–272.

Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. *Dbpedia: A nucleus for a web of open data*. Springer.

Collin F Baker, Charles J Fillmore, and John B Lowe. 1998. The berkeley framenet project. In *Proceedings of the 17th international conference on Computational linguistics-Volume 1*, pages 86–90. Association for Computational Linguistics.

Niranjan Balasubramanian, Stephen Soderland, Oren Etzioni Mausam, and Oren Etzioni. 2013. Generating coherent event schemas at scale. In *EMNLP*, pages 1721–1731.

Dana H Ballard and Patrick J Hayes. 1986. *Parallel logical inference and energy minimization*. Department of Computer Science, University of Rochester.

Ann M Bisantz and Kim J Vicente. 1994. Making the abstraction hierarchy concrete. *International Journal of human-computer studies*, 40(1):83–117.

Daniel G Bobrow and Terry Winograd. 1977. An overview of krl, a knowledge representation language. *Cognitive science*, 1(1):3–46.

Daniel G Bobrow. 1964. Natural language input for a computer problem solving system.

Margaret A Boden. 1996. The philosophy of artificial life.

Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, pages 1247–1250. ACM.

Alexander Borgida, Ronald J Brachman, Deborah L McGuinness, and Lori Alperin Resnick. 1989. Classic: A structural data model for objects. In *ACM Sigmod record*, volume 18, pages 58–67. ACM.

Rodney A Brooks. 1990. Elephants don't play chess. *Robotics and autonomous systems*, 6(1):3–15.

Rodney A Brooks. 1991. Intelligence without representation. *Artificial intelligence*, 47(1):139–159.

Bruce G Buchanan and Edward A Feigenbaum. 1978. Dendral and meta-dendral: Their applications dimension. *Artificial intelligence*, 11(1):5–24.

Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R Hruschka Jr, and Tom M Mitchell. 2010. Toward an architecture for never-ending language learning. In *AAAI*, volume 5, page 3.

Eugene Charniak. 1972. Toward a model of children's story comprehension.

Peter Cheeseman. 1985. In defense of probability. In *IJCAI*, pages 1002–1009.

Noam Chomsky. 1988. *Language and problems of knowledge: The Managua lectures*, volume 16. MIT press.

Peter Clark. 2015. Elementary School Science and Math Tests as a Driver for AI:Take the Aristo Challenge! In *Proceedings of IAAI*.

R.T. Cox. 1961. *Algebra of Probable Inference*. Johns Hopkins University Press.

Ido Dagan, Bill Dolan, Bernardo Magnini, and Dan Roth. 2010. Recognizing textual entailment: Rational, evaluation and approaches–erratum. *Natural Language Engineering*, 16(01):105–105.

Adnan Darwiche and Moisés Goldszmidt. 1994. On the relation between kappa calculus and probabilistic reasoning. In *Proceedings of the Tenth international conference on Uncertainty in artificial intelligence*, pages 145–153. Morgan Kaufmann Publishers Inc.

Gerald Francis Dejong, II. 1979. *Skimming Stories in Real Time: An Experiment in Integrated Understanding.* Ph.D. thesis, New Haven, CT, USA. AAI7925621.

Hubert L Dreyfus. 1992. *What computers still can't do: a critique of artificial reason*. MIT press.

Jeffrey L Elman. 1991. Distributed representations, simple recurrent networks, and grammatical structure. *Machine learning*, 7(2-3):195–225.

Jeffrey L Elman. 1996. *Rethinking innateness: A connectionist perspective on development*, volume 10. MIT press.

Oren Etzioni, Michael Cafarella, Doug Downey, Stanley Kok, Ana-Maria Popescu, Tal Shaked, Stephen Soderland, Daniel S Weld, and Alexander Yates. 2004. Web-scale information extraction in knowitall:(preliminary results). In *Proceedings of the 13th international conference on World Wide Web*, pages 100–110. ACM.

Oren Etzioni. 1993. Intelligence without robots: A reply to brooks. *AI Magazine*, 14(4):7.

Christiane Fellbaum. 1998. *WordNet*. Wiley Online Library.

Richard Fikes and Tom Kehler. 1985. The role of frame-based representation in reasoning. *Communications of the ACM*, 28(9):904–920.

Charles J Fillmore. 1977. Scenes-and-frames semantics. *Linguistic structures processing*, 59:55–88.

Jerry A Fodor and Zenon W Pylyshyn. 1988. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1):3–71.

Jerry A Fodor. 1983. *The modularity of mind: An essay on faculty psychology*. MIT press.

Ken-Ichi Funahashi. 1989. On the approximate realization of continuous mappings by neural networks. *Neural networks*, 2(3):183–192.

Matt Gardner, Partha Pratim Talukdar, Bryan Kisiel, and Tom M Mitchell. 2013. Improving learning and inference in a large knowledge-base using latent syntactic cues. In *EMNLP*, pages 833–838.

E Goldstein. 2014. *Cognitive psychology: Connecting mind, research and everyday experience*. Cengage Learning.

Charles M Gray and Wolf Singer. 1989. Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proceedings of the National Academy of Sciences*, 86(5):1698–1702.

Frederick Hayes-Roth, Donald Waterman, and Douglas Lenat. 1984. Building expert systems.

Patrick J. Hayes. 1995. Computation & intelligence. chapter The Second Naive Physics Manifesto, pages 567–585. American Association for Artificial Intelligence, Menlo Park, CA, USA.

Steffen Hölldobler, Yvonne Kalinke, Fg Wissensverarbeitung Ki, et al. 1991. Towards a new massively parallel computational model for logic programming. In *In ECAI94 workshop on Combining Symbolic and Connectioninst Processing*. Citeseer.

Alan Hjek. 2012. Interpretations of probability. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Winter 2012 edition.

Michael E Janzen and Kim J Vicente. 1997. Attention allocation within the abstraction hierarchy. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 41, pages 274–278. SAGE Publications.

Edwin T Jaynes. 2003. *Probability theory: the logic of science*. Cambridge university press.

Addie Johnson and Robert W Proctor. 2004. *Attention: Theory and practice*. Sage Publications.

David Kirsh. 1991. Today the earwig, tomorrow man? *Artificial intelligence*, 47(1):161–184.

Ni Lao, Tom Mitchell, and William W Cohen. 2011. Random walk inference and learning in a large scale knowledge base. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 529–539. Association for Computational Linguistics.

Steffen L Lauritzen. 1996. *Graphical models*. Oxford University Press.

Douglas B Lenat. 1995. Cyc: A large-scale investment in knowledge infrastructure. *Communications of the ACM*, 38(11):33–38.

Hector J Levesque, Ernest Davis, and Leora Morgenstern. 2011. The winograd schema challenge. In *AAAI Spring Symposium: Logical Formalizations of Commonsense Reasoning*.

Hugo Liu and Push Singh. 2004. Conceptneta practical commonsense reasoning tool-kit. *BT technology journal*, 22(4):211–226.

Pattie Maes. 1990. Situated agents can have goals. *Robotics and autonomous systems*, 6(1):49–70.

Mausam, Michael Schmitz, Robert Bart, Stephen Soderland, and Oren Etzioni. 2012. Open language learning for information extraction. In *Proceedings of Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*.

John McCarthy and Patrick Hayes. 1968. *Some philosophical problems from the standpoint of artificial intelligence*. Stanford University USA.

John McCarthy and Vladimir Lifschitz. 1990. *Formalizing common sense: papers*, volume 5. Intellect Books.

John McCarthy. 1963. *Programs with common sense*. Defense Technical Information Center.

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

George A Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.

Marvin Minsky and Seymour Papert. 1969. Perceptron: an introduction to computational geometry. *The MIT Press, Cambridge, expanded edition*, 19:88.

Marvin Minsky. 1974. A framework for representing knowledge.

Marvin Minsky. 1988. *Society of mind*. Simon and Schuster.

T. Mitchell, W. Cohen, E. Hruscha, P. Talukdar, J. Betteridge, A. Carlson, B. Dalvi, M. Gardner, B. Kisiel, J. Krishnamurthy, N. Lao, K. Mazaitis, T. Mohammad, N. Nakashole, E. Platanios, A. Ritter, M. Samadi, B. Settles, R. Wang, D. Wijaya, A. Gupta, X. Chen, A. Saparov, M. Greaves, and J. Welling. 2015. In *AAAI*.

Erik T Mueller. 1998. *Natural language processing with ThoughtTreasure*. Citeseer.

Roberto Navigli and Simone Paolo Ponzetto. 2010. Babelnet: Building a very large multilingual semantic network. In *Proceedings of the 48th annual meeting of the association for computational linguistics*, pages 216–225. Association for Computational Linguistics.

Allen Newell and Herbert Alexander Simon. 1961. Gps, a program that simulates human thought. Technical report, DTIC Document.

Allen Newell and Herbert A Simon. 1976. Computer science as empirical inquiry: Symbols and search. *Communications of the ACM*, 19(3):113–126.

Nils J Nilsson. 1986. Probabilistic logic. *Artificial intelligence*, 28(1):71–87.

Judea Pearl. 1986. Fusion, propagation, and structuring in belief networks. *Artificial intelligence*, 29(3):241–288.

Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. *Proceedings of the Empiricial Methods in Natural Language Processing (EMNLP 2014)*, 12.

Gadi Pinkas. 1991. Symmetric neural networks and propositional logic satisfiability. *Neural Computation*, 3(2):282–291.

Steven Pinker and Paul Bloom. 1990. Natural language and natural selection. *Behavioral and brain sciences*, 13(04):707–727.

Jens Rasmussen. 1985. The role of hierarchical knowledge representation in decisionmaking and system management. *Systems, Man and Cybernetics, IEEE Transactions on*, (2):234–243.

Matthew Richardson, Christopher JC Burges, and Erin Renshaw. 2013. Mctest: A challenge dataset for the open-domain machine comprehension of text. In *EMNLP*, pages 193–203.

R Bruce Roberts and Ira P Goldstein. 1977. The frl manual. Technical report, DTIC Document.

Frank Rosenblatt. 1958. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6):386.

David E Rumelhart and James L McClelland. 1985. On learning the past tenses of english verbs.

David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. 1988a. Learning representations by back-propagating errors. *Cognitive modeling*, 5.

David E Rumelhart, James L McClelland, PDP Research Group, et al. 1988b. *Parallel distributed processing*, volume 1. IEEE.

Klaus Schild. 1991. *A correspondence theory for terminological logics: Preliminary report*.

Manfred Schmidt-Schauß and Gert Smolka. 1991. Attributive concept descriptions with complements. *Artificial intelligence*, 48(1):1–26.

John R Searle. 1980. Minds, brains, and programs. *Behavioral and brain sciences*, 3(03):417–424.

Glenn Shafer. 1976. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton.

Lokendra Shastri and Venkat Ajjanagadde. 1993. From simple associations to systematic reasoning: A connectionist representation of rules, variables and dynamic bindings using temporal synchrony. *Behavioral and brain sciences*, 16(03):417–451.

Push Singh, Thomas Lin, Erik T Mueller, Grace Lim, Travell Perkins, and Wan Li Zhu. 2002. Open mind common sense: Knowledge acquisition from the general public. In *On the Move to Meaningful Internet Systems 2002: CoopIS, DOA, and ODBASE*, pages 1223–1237. Springer.

Paul Smolensky. 1988. The constituent structure of connectionist mental states: A reply to fodor and pylyshyn. *The Southern Journal of Philosophy*, 26(S1):137–161.

Paul Smolensky. 1990. Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial intelligence*, 46(1):159–216.

Robert Speer and Catherine Havasi. 2012. Representing general relational knowledge in conceptnet 5. In *LREC*, pages 3679–3686.

Wolfgang Spohn. 1988. *Ordinal conditional functions: A dynamic theory of epistemic states*. Springer.

Matthias Steup. 2014. Epistemology. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Spring 2014 edition.

Fabian M Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. Yago: a core of semantic knowledge. In *Proceedings of the 16th international conference on World Wide Web*, pages 697–706. ACM.

Lucy A Suchman. 1987. *Plans and situated actions: the problem of human-machine communication*. Cambridge university press.

David S Touretzky and Geoffrey E Hinton. 1985. Symbols among the neurons: Details of a connectionist inference architecture. In *IJCAI*, volume 85, pages 238–243.

Geoffrey G Towell and Jude W Shavlik. 1993. Extracting refined rules from knowledge-based neural networks. *Machine learning*, 13(1):71–101.

Anne M Treisman and Garry Gelade. 1980. A feature-integration theory of attention. *Cognitive psychology*, 12(1):97–136.

Amos Tversky and Daniel Kahneman. 1974. Judgment under uncertainty: Heuristics and biases. *science*, 185(4157):1124–1131.

Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Communications of the ACM*, 57(10):78–85.

Jason Weston, Antoine Bordes, Sumit Chopra, and Tomas Mikolov. 2015. Towards ai-complete question answering: A set of prerequisite toy tasks. *arXiv preprint arXiv:1502.05698*.

Terry Winograd. 1971. Procedures as a representation for data in a computer program for understanding natural language. Technical report, DTIC Document.

Wentao Wu, Hongsong Li, Haixun Wang, and Kenny Q Zhu. 2012. Probase: A probabilistic taxonomy for text understanding. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*, pages 481–492. ACM.

Lotfi A Zadeh. 1973. Outline of a new approach to the analysis of complex systems and decision processes. *Systems, Man and Cybernetics, IEEE Transactions on*, (1):28–44.

LA Zedeh. 1989. Knowledge representation in fuzzy logic. *Knowledge and Data Engineering, IEEE Transactions on*, 1(1):89–100.